Issues in the Design of a Multistep Code

Fred T. Krogh[1,*]

[1] *Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109, U.S.A.*

This paper summarizes things the author considers important from his past work, and presents some new ideas for use in the implementation of multistep methods. An effort has been made to include material of interest to those of a theoretical bent, as well as to those whose primary interest is in the implementation of methods.

1. Introduction

This paper is a summary of my views on the integration of differential equations. I hope the following groups will find something of interest in what follows.

For those interested in the implementation of multistep methods, a new approach to starting should be of particular interest. Topics in my published work mentioned here include the direct integration of higher order equations, and a test for noise in the computation, as I feel these items have not had the attention they deserve. Ideas not previously published on order selection and the treatment of discontinuities and constraints are outlined. A new code I plan to write which will support the integration of stiff and differential algebraic systems, delay equations, and of course the efficient integration of nonstiff differential equations is discussed briefly.

For those interested in one step methods, this paper offers the chance to spy on some of the competition. It is generally agreed that multistep methods work

---

best if derivative evaluations are expensive or if there are complex output requirements. Research on one step methods has both improved their efficiency, and provided for more flexibility in output. Still, it will be difficult to support with a one step code the wide class of problems discussed here to be supported with a multistep code. And if one can write such a one step code it is liable to suffer on integration overhead and code complexity to the point that it will have lost the big advantages one step methods possess.

For those interested in theory, the next section is intended for you. In addition, a few comments appear in bold face to attract your attention. I am truly in awe at the theoretical results obtained concerning the numerical solution of ode's. But for an algorithm developer, these results have for the most part been irrelevant. I have never had occasion to use your work, unless it was a result you obtained in the process of writing a code. When needing a result of a theoretical nature while developing a code, I have been on my own. Based on a sample of two other code developers, this is probably the typical situation. As one whose interests and talents are algorithmic rather than mathematical, I would like to see more help from those with the opposite inclinations.

Finally, I would like to comment on approaches to developing algorithms for a code. Not counting the time spent coding or debugging, I have spent approximately 30% of my time working on theory, 5% on the shotgun approach of adjusting an algorithm based on results for a big set of test problems, and 65% of my time with what I call the intelligent observer approach. This approach involves solving a very simple problem for which one knows the desired behavior based on theoretical considerations. Then watching closely what the algorithm does concerning such things as estimating errors and selecting order and stepsize. And asking myself over and over again, what would I do here knowing what I do about the problem, and about the non-problem specific information which could be made available to the program. When unexpected behavior is encountered, I ask, "Why did that happen?" Because of these efforts, my programs can do a better job of selecting order and stepsize than I can! This approach has also worked well in developing algorithms for nonlinear least squares and quadrature.

2. Rigorous selection of stepsize and order is not a piece of cake

Perhaps the most important area where theory could be helpful is in the selection of stepsize and order. But for multistep methods reality is more complicated than the theory constructed to address these issues. When solving quadrature problems, existing theory adequately addresses the issues. But with quadrature problems the usual problems with instability do not arise.

It is useful to consider the simple test problem $y' = \lambda y$ for constant complex λ. Although this simple problem does not capture all the complexity of real problems, it seems to capture enough to be useful; and it illustrates the problems

with the theoretical approaches. I assume the stepsize, h, is constant, even while talking out of the other side of my mouth about changing it. As an algorist, and as one who has spent much time looking at difference tables verifying that this assumption seems useful, I have two excuses for such behavior. Perhaps you mathematicians could do better. (That is, either in behavior or excuses.)

For the case above, multistep methods are characterized by having a principal root, r_p, and $c_p r_p^n$ approximates the solution with order $q+1$, where q is the order of the last difference used and c_p is a constant. Unfortunately, they also have extraneous roots which are introduced as a result of the difference equation having a higher order than the differential equation. For all practical methods I know of, the largest of these extraneous roots, r_e, near the stability boundary for a method has a negative real part. Accuracy considerations on the other hand, limit the argument of r_p to $< 60°$. This corresponds to only 6 samples per period; 12 samples ($30°$) is probably more realistic.

Since the k^{th} backward difference of r^n is $r^n (1 - 1/r)^k$, both error estimation and order selection depend on expressions approximated by

$$c_p r_p^n (1 - 1/r_p)^k + c_e r_e^n (1 - 1/r_e)^k.$$

Because of their arguments, the contribution of r_p to the k^{th} difference decrease with increasing k, while the opposite is true for the contribution of r_e. Since, $|c_p| < |c_e|$ (if there is any relative accuracy at all), the magnitude of the differences decreases at first, but ultimately increases. A variable order method selects the order very near this turning point. Thus the quantities used for decision making depend on both r_p and r_e, although the standard theory only applies to r_p. Although $(1 - 1/r_p)^k$ is proportional to h^k, $(1 - 1/r_e)^k$ will actually decrease with increasing h since $Re(r_e) < 0$, and $|r_e|$ increases with increasing h. Ultimately, and quite quickly, this reduction in $|1 - 1/r_e|^k$ is more than compensated for by the increase in $|r_e|^n$. Of course, if the lack of convergence in the differences is due to round off error, an increase in h will lead to a smaller than expected increase in the error, particularly if this larger stepsize results in an increase in the order. Except for the comment on round off error, all said above about increasing h, also applies to increasing the order. That is, an increase in the order tends to make $|r_e|$ larger.

Of course changing stepsize or order will in effect change the value of c_e as well as the value of r_e. More important, the complex numbers can't be examined directly. In a system of real equations, the complex quantities all get combined in such a way that the result is real. Imagine for example trying to determine if $|\rho^k e^{ik\theta}|$ is an increasing function of k given only $\rho^k \cos k\theta$, for a few values of k. More than one r_p and/or r_e complicates matters even more. That is, if one is still interested in trying to treat this problem rigorously.

I have in fact observed differences not increasing as rapidly as expected after a stepsize increase. And sometimes a questionable order increase is followed by

a clearly undesirable one. The cautious will claim this shows one should not allow stepsize or order to change too frequently. But with such prudence, one pays a price in performance when a problem should have just such rapid changes. One is probably better off overall to make small mistakes in stepsize or order selection, and then to correct mistakes as quickly as possible.

In the past I have selected order independently for different equations in the belief that this would pay big dividends when a system of equations involves equations with very different characteristics. As mentioned later, in my new integrator users will be encouraged to partition their equations. Usually users can do a decent job of this, and as a result I plan to use the same order for all of the equations in a partition. Of course, this does not change the fact that **error estimates do not behave in the same way as actual local errors.**

3. Features from the Past

Table 1 lists features in my variable order Adams code, DIVA [13]. This code is essentially the same as a code finished in 1975, except for minor changes in the algorithms and the addition of features requested since then. The main ideas in the algorithms used are in [9]. A more accessible reference is [17]. However, this latter reference does not include details for higher order equations or for implementing the backward differentiation formulas, and does not give the most efficient algorithm for computing interpolation coefficients.

Table 1
Features in DIVA

G-Stops -- Locating zeros of arbitrary functions of the solution.
Direct integration of second order systems.
Saving the solution of later use.
Internal diagnostic print, and print for help with debugging.
Support for integration of variational equations.
Flexibility in specifying output points.
Reverse communication is supported, but not required.
Flexible error control.
User can control stepsize and order.
Diagnoses unreasonable accuracy requests.
Low overhead compared to other variable order Adams methods.
Efficient integration over weak discontinuities.

Jackson [4], examines a variety of methods proposed for computing with divided differences or variants thereof, and concludes the algorithms in [9] are best. Jackson and Sedgwick [5], examine the tradeoffs between using the

Lagrange, the Nordsiek, or the form used in [9] for the interpolating polynomial, and conclude the last is to be preferred. In this paper they suggest using the term "scaled" divided differences for what I have called "modified" divided differences. I picked up the less descriptive term "modified" from a paper by Blanch [1], which in fact uses a different definition. I plan to use the more descriptive term "scaled" in the future.

All of the features in Table 1 have been used; some of them have been added because of user requests. I would like to comment on a few of these features.

Allowing the direct integration of second order systems costs very little in the code. The procedure for computing integration coefficients for a first order system, computes in the process, coefficients for integrating higher order systems. And of course one only needs to maintain half the number of difference tables, so the integration overhead per equation is cut by a factor of almost 2. In the case of implicit methods, one also saves on the linear algebra cost. There are important classes of problems [11], where the direct integration of a second order system requires about half the number of function evaluations required for integrating the equivalent first order system. It should be noted that if the first derivative of the solution is present in the expressions for the second derivatives, it is possible for the direct integration to require more function evaluations. Even in this case, the reduction in integration overhead is likely to make up for the small number of extra function evaluations required in the worst case. The error in second order systems is controlled by allowing the user to specify an error tolerance only on the first derivative of the solution. This has worked well in practice, and eliminates the possibility of inconsistent accuracy requests on position and velocity. **However, I know of no theoretical results indicating the best strategy.**

Diagnostic print is intended to aid in the identification of problems in the algorithm or errors in the user's problem formulation.

Debug print is intended primarily to help in tracking down user coding errors. Instead of telling the user (as in the distant past) to print this and that, to track down problems, the user is told to turn on this kind of output, at that level. From experience with the former approach I know one gets poorly labeled output, badly organized on the page, with not quite enough digits on certain key variables, and some desired data missing!

The support for variational equations uses a method of the form $PE_1C_1E_2C_2E$, where subscripts indicate the action is to apply only to the first or second set of equations, and P, E, and C are used to denote predicting, evaluating the derivatives and correcting respectively. Letters without subscripts imply the action is to apply to both sets. Values of the solution for variables in the second set have no influence on the derivatives of those in the first set. Thus in the case of variational equations, the partials matrix need only be computed once per step, such values being computed from the corrected values for the state. This feature is quite easy to implement, and has a big payoff for some problems.

In [9], a method is described for detecting computational noise. This test does not make assumptions about the precision with which derivatives are computed, and as a result has found the following kinds of errors:

1. User thought all of his code was in double precision, when in fact certain computations were done in single.
2. User thought computing attraction of sun and earth in double precision, while computing the rest in single precision (to save time and space) would be sufficient to get the desired accuracy. It usually was -- unless the spacecraft got a little too close to the moon.
3. User did not realize computing with time based back at the time of Caesar caused cancellation which made the desired accuracy impossible to obtain.

Tests based on the precision of the floating point arithmetic have no chance of catching such errors, and these problems are of a nature where they could easily go undetected. The test is based on examining the rate with which differences get smaller, together with the place where the decrease ends. (If r_p is so close to 1 that differences start decreasing rapidly, then $h\lambda$ must be close to 0, and thus r_e must be small. If r_e is small then lack of convergence must be due to noise.) As given in [9], this test can give a false positive result, and so has been modified to form a second difference of derivative values for very closely spaced values of the independent variable. The diagnostic is only given when this difference is sufficiently large. These two extra derivative evaluations are requested when the test in [9] indicates problems, which happens very rarely. As part of the testing process, the step is made easier to increase when it looks like noise might be limiting accuracy, and thus the overall effect of the test is to make the code more efficient when noise is (nearly) limiting precision. **Perhaps some test for noise could be placed on some kind of theoretical foundation?** In the past I have not tried to estimate Lipschitz constants which in turn could be used to estimate a bound on the eigenvalues. Such estimates might prove useful both here and in selecting stepsize and/or order.

When saving the solution, restarts require more storage to save the solution because the low order of the difference tables requires saving information more frequently. This became a problem when integrating trajectories for earth satellites which pass in and out of the Earth's shadow. It is easy to imagine such concerns also arising in connection with integrating delay equations. Simply ignoring discontinuities is very unreliable if one is using my code, while the infamous John Butcher assures me that the one step code STRIDE is reliable in this case. Thus, either the unreliability is common to multistep methods, rooted in attempting to use invalid past information, or is due to DIVA (for good reasons) not allowing the order to drop by more than one on a single step. But even if ignoring discontinuities were reliable, doing so is still quite inefficient.

I have investigated two approaches to integrating over weak discontinuities.

Both require the user to identify the amount and the point of discontinuity, and both assume information in the differences prior to the discontinuity are of some use in predicting the differences after the discontinuity. The first approach involves doing a Picard like starting procedure using the new definition for the derivatives prior to the point of the discontinuity, and starting with the same values for all of the differences except for the 0^{th}. This procedure worked well for small discontinuities, but performance degraded fairly rapidly with increasing size of the discontinuity. More important, there are cases where it is not possible to extend the new derivative evaluation to points prior to the discontinuity. Thus my users had to do something quite special as the spacecraft moved from full sun or shadow into partial shadow. (There is no problem extending the derivatives when moving from partial shadow, as full sun or shadow is easy to model; extending partial shadow is impossible.)

The second approach is more in the spirit of variable order methods. Upon being informed of the discontinuity and its size, DIVA reduces the integration order to the point where the largest order difference is larger than the size of the discontinuity, and adjusts the stepsize accordingly. A little special logic is required to get the solution on the first step past the discontinuity. Then for a number of steps depending on the integration order, DIVA is extra careful in terms of being a little more conservative in estimating errors, and being willing to reduce the order more quickly. Overall the two approaches were about equally good, although this last approach is distinctly better for large discontinuities. Using the idea described later for starting the integration (i.e. removing points just after the discontinuity before the point at the discontinuity) should make this second approach work even better.

Thus frequent discontinuities do not necessarily give one step methods an advantage as has been claimed. And if the discontinuities are frequent enough, a variable order code running at first order would be impossible to beat!

4. Features for the Future

Table 2 lists the features I plan to put into a new integrator. More details can be found in [14]. Other obligations, will probably delay the start of programming till well into 1994.

Topics from Table 2 are discussed below. Constraints on the size of this paper limit the detail with which these topics can be discussed.

Various aspects of how I plan to use the BDF's for stiff equations can be found in [15], [16], and [19]. Unlike most authors using the BDF's, I do not regard the method as one where one iterates until getting a solution to the corrector. Rather, the method is semi-implicit using some fixed number of iterations. The method uses the true variable coefficient value in computing the residual for the equation to be solved. It is not true, as is frequently thought, that such a method must factor the iteration matrix frequently. Results given in

[19] certainly suggest that this general approach is promising.

Table 2
Features Planned for a New Integrator

Everything in Table 1
Stiff Equations -- plan to used backward differentiation formulas, BDF's.
Differential Algebraic Systems, DAE's (index $\leq$ 1).
Delay Equations.
Systems subject to constraints.
Support for changing the independent variable.
Support for using extra derivatives.
An improved starting procedure.
A simplified error control.
Provision for partitioning equations to use different methods.
Automatic selection of method, Adams or BDF.
Provision for carrying additional precision in independent and dependent variables.
Continuous order selection?
Provision to get a solution depending continuously on the input parameters?
Provision for different stepsizes on different equations?
A new user interface

The restriction to index 1 on the DAE's is not expected to unduly constrain the range of applications, while greatly simplifying the code development. For higher index problems, the user will need to reduce the index by differentiating (some of) the equations, and solve a system consisting of differentiated equations together with equations which have not been differentiated, subject to constraints based on the differentiated equations prior to being differentiated. The feature for imposing constraints on the solution is used to satisfy these constraints.

At the request of a user, a modification was made to my current integrator to solve a delay equation. This problem involved delays from tidal forces propagating through the moon, with the state to the left of the initial point to be determined in such a way as to give a smooth solution. This initial state was determined by initially assuming the values at the delayed time were the same as at the current values, and doing a Picard like iteration to get this initial state. Since many physical processes are characterized by smooth solutions, it is reasonable to expect this kind of problem is common. It appears one step methods would have difficulty treating such problems.

I first suggested modifying the numerical solution of a differential equation to comply with constraints in [8], and a number of others have independently come up with this idea. The basic idea is to project onto the constraint just after correcting and just before computing the last derivative evaluation of the step. The projection involves one iteration of solving an underdetermined system in

the linearized constraints. Thus, the cost is quite low. With just a single constraint, the cost can be made even lower by moving onto the constraint in the direction of the difference between the predicted and corrected values. **(At least this worked very well for the one problem on which it was tried!)** This latter procedure requires computing the value of the constraint after predicting and correcting, instead of requiring the value of the constraint and it's partial derivatives after correcting. This technique was tried on a two-body problem with eccentricity .5 over 21870 (lots!) of periods, constraining the energy to be constant. Without the constraint errors grow quadratically with time; with it, the error grows linearly. In comparison with the best results by Calvo and Sanz-Serna [2] for symplectic integrators, DIVA without constraints took less than half as many function values for a given accuracy, but could not get errors below 10^{-3}, where [2] gives results almost down to an accuracy of 10^{-4}. With constraints, DIVA with an error 1/100 of those reported in [2], required less than 1/10 as many function evaluations. Though this problem does not show symplectic integrators at their best (because of the mild eccentricity), it certainly suggests that the use of constraints with an existing integrator, offers better performance than symplectic integrators when such a procedure is feasible. For those interested in making comparisons, results with the constraint are given for DIVA on this problem in the form (function evaluations, largest error through 21870 periods): $(4372756, 1.22\times10^{-3})$, $(5213486, 8.38\times10^{-5})$, $(6147913, 3.54\times10^{-6})$, $(7106722, 1.10\times10^{-6})$, and $(8047023, 1.00\times10^{-7})$. These results were obtained with input tolerances of 10^{-k}, for k = 7, 8, ..., 11, and $(x_0, x_0', y_0, y_0') = (.5, 0, 0, 3.)$.

There are many problems where a change in independent variable can make a significant difference in the efficiency of an integration. An example is given in [10], where a factor of 2 improvement is realized. For problems with very rapid changes in the solution, people at TRW were using arc length as an independent variable at least 25 years ago. But if a user wants to make such a change in variable, getting output at the desired points can become much messier to set up. I would like to encourage the user to make changes in the independent variable when it is desirable, by setting up the code to use the G-Stop feature on the solution for the original independent variable to get output at the points where it is desired. Thus the user has little to do other than defining $dt/d\tau$, where t is the old and τ the new independent variable. Note that the iteration to find the G-Stop for a particular value of t, only requires the interpolations be done for that variable. The remaining interpolations can be done after the desired value for τ is obtained.

I experimented some with using extra derivatives in [6]. If one has an integrator set up to solve arbitrary order differential equations directly, this feature is quite easy to add. One treats the system as if it had order d+e where d is the original order and e is the number of extra derivatives. Except, the user provided values for derivatives from d to e-1 are used instead of the values ordinarily computed by the integrator for an equation of order d+e. Long ago

I felt the extra complexity was not worth the cost and thus integrators I have written since have not had this feature. The expected development of convenient software for automatic differentiation, and the chance that the use of extra derivatives will frequently enable effective use of PEC mode when PECE would otherwise be preferred, has prodded me into including this feature. **The error control here will be based on the d−1st derivative, not on the (d + e)−1st derivative.**

The new approach to starting arose as a result of work on integrating past discontinuities. The failure of one "great" idea could only be explained by the extreme growth of rounding errors in the presence of a rapidly increasing stepsize. Clearly the same problem arises when starting an integration with a first order method. The problem is easy to see when working with scaled divided differences. But, please note that the use of differences is not the source of the difficulty. Assume k steps in the starting process with the order increased by 1 on each step. If the order is then to be held constant (or perhaps even reduced), the usual procedure is to discard information from the most distant point, i.e. the initial point. But observe that the initial point is the most accurate information available, and that there is a great deal of cancellation in forming the difference between the initial point and the point following. If the point just after the starting point is removed, the shortest distance between two points is maximized, thus reducing as much as possible the growth of roundoff error.

The easiest way to implement this starting procedure is probably to keep the initial point as the base point of the integration until it is discarded. Thus after getting values at t_n (the current value of the independent variable) they are inserted into the difference tables between t_{n-1} and t_0. The process for updating the differences is a little different, but, except for the stepsize history, the procedure for computing integration coefficients remains the same. This process has the advantage of integrating from t_0 to t_n, instead of from the t_{n-1} to t_n. Because of the way errors in the interpolating polynomial cancel, this is just as accurate, and in fact will be slightly more accurate if some of the values discarded early are not quite as accurate as they should be. During this process one need only compute one derivative value per step. When the difference between the t_0 and the t_k closest to t_0, is nearly as big as $|t_n - t_{n-1}|$, it is time to modify the difference tables by discarding the data at t_0. This is also easy to do. If a history is being saved for later interpolation of the solution, it should be saved just before discarding the data for t_0. Note that the usual starting procedure requires more storage for saving the solution, since it must save low order closely spaced difference tables near the starting point.

In the past I have tried to provide a way for the user to control the accuracy by means of a single parameter. **Ideally there would be a theoretical result such as: given the information accessible to a code, this is the best error control to use.** Lacking such a result, an error relative to the change in the solution over a single step looks good to me. To avoid problems with zero crossing an exponential averaging technique is used. Thus for some α ($\approx 3/4$), use an error

tolerance relative to $\alpha\times$(last used) + $(1-\alpha)\times$(current relative change). Unfortunately, the starting stepsize at first order is abnormally small. Here one would like to use the expected value for the stepsize assuming a more realistic value for the integration order. This little problem has kept me from implementing this procedure. I'm hoping the new starting procedure described above will make this type of simplified error control feasible. The goal is not to do the best possible thing, but to do at least as well as a user is likely to do when he is uncertain about the nature of the error growth and the size of the solutions in his problem.

My first paper on stiff equations, [7], advocated partitioning the equations to reduce the size of the linear systems involved in getting a solution. The simple minded approach suggested there for doing automatic partitioning is probably not workable. A rough outline for a more sophisticated approach is given in [12]. But there are many cases where the person with the problem knows how equations should be grouped so that it is appropriate to integrate all of the equations in the group with the same method. Thus, for example, control equations for keeping a spacecraft oriented belong in a different group than those which model the external forces on the spacecraft. I believe software for solving difficult problems should use all available help. The user interface described in [14] attempts to encourage a user in this direction.

I have users who are interested in getting solutions pushing the accuracy available with IEEE double precision format. One can pick up a little accuracy by maintaining the independent and dependent variables in extra precision. I plan to offer such an option. This will require hand coding extra precision addition only, which is quite simple. The basic idea goes back to Gill, [3]. (One of my first programs ($\approx$ 1960), and the only Runge-Kutta method I've written, used this method coded in absolute hex on an ALWAC III-E. Assembly language, if even available then, didn't allow all the tricks that were necessary!)

When selecting the integration order, particularly at high order, it is easy to be at an order where everything is nicely stable, and getting further outside the region of relative stability than is desirable with the order increased by 1. This suggests one might use a method of "real" order. Let q be such a real order, and let k be the integer part of q. The method being proposed is simply the usual method of order k, with (q-k) times the first neglected term in both the predictor and the corrector added to the usual formulas. Clearly such a method, although of order k, will have properties between those for the method of order k and the method of order k+1. Another reason for considering such an algorithm is to open the possibility of getting a solution that varies continuously as a function of the initial conditions or some parameter appearing in the differential equation. This feature is useful when solving a boundary value problem when one does not want to integrate variational equations to get partials. Without this feature, one cannot get decent values for partial derivatives by using difference approximations. I have only recently realized that

getting this property when integrating a stiff equation will probably require computing the Jacobian matrix and factoring the iteration matrix on every step. One would probably be better off integrating variational equations. Thus, some of my enthusiasm for using real values of the integration order has been dampened. However, the idea seems of some interest, and is used in the next section of this paper in investigating "optimal" formulas.

Multirate methods are those which allow different stepsizes to be used on different equations. This option was originally specifically excluded because of the code complexity involved in including it. Since then, two parties have been interested in this feature, and I have kidded myself into believing that perhaps the implementation isn't too messy. I believe a key to making these methods work is to use some kind of averaging to feed the contribution from the small stepsize components into the derivatives computed for the large stepsize components.

One aspect of the user interface is crucial. Whatever approach to a user interface you use, try to make it extensible. I have modified an existing code over a period of over 15 years, in the process adding significant features not in the original design. Past use of the program was invalidated once to provide better portability of saved solutions created on different computer systems. (It's hard to think of everything!) Since features could be added without invalidating existing code, one version of the program was maintained, and everybody using the code could make use of the latest version. Although the design is different, the new integrator keeps this important characteristic. In particular, no matter what kind of method used to solve the corrector equation when using an implicit method, the same core code will be used. If one wants to use the latest sparse solver, it should just be a matter of adjusting the user interface slightly for that code, and using it.

5. What is an Optimal Formula?

Long ago there were efforts to improve the order or the accuracy of formulas, without any consideration of stability. Then, there were efforts to improve stability with only minimal consideration of accuracy. Of course, improving one usually means damaging the other. As one wants more accuracy, it is best to use more accurate formulas, sacrificing stability to get the accuracy, and vice versa. Variable order methods, provide a means to pick the nearly best compromise automatically. But this leaves open the question of just which variable order formulas to use. Skeel [18], has shown it is possible to have a great deal of freedom in the choice of formulas without sacrificing storage. The best choice may indeed lie in such formulas. I have wondered for some time about doing something less general, but significantly easier to implement.

The corrector formula for the Adams methods applied to first order systems,

can be written as: $y_{n+1} = p_{n+1} + h\,c\,(f(t_{n+1}, p_{n+1}) - p'_{n+1})$, where y and p denote corrected and predicted values, c is a constant which depends on the method, f is the derivative function, and p'_{n+1} is the value predicted for the derivative by extrapolating the polynomial interpolating past derivative values at the end of the last step. Clearly there is freedom in choosing c, with little impact on the complexity of the method. In the past I have selected c so that the corrector formula has order one greater than that of the predictor, primarily because using a value for c giving the Adams Moulton formula with the same order as the predictor is not as stable for $y' = -y$ for predictor orders ≤ 11.

The issue addressed here is how to decide one value of c is better than another, when a change in c makes the formula more accurate and less stable or vice versa. Since order selection ultimately depends on numbers actually used in the computation, define the method to be stable if $|r_c(h\lambda)| \le |r_p(h\lambda)|$. That is, use relative stability, relative to the principal root rather than to the true solution. If this condition is violated, then ultimately a variable order method will drop the order. The key to choosing the optimal value of c, is using a real valued order as introduced earlier.

With a predictor "order" of $q = k + \rho$, $0 \le \rho < 1$, and free parameter α, the method in backward difference form can be written

$$p_{n+1} = y_n + h\sum_{j=0}^{k-1} \gamma_j \nabla^j y'_n + \rho\,\gamma_k \nabla^k y'_n$$

$$p'_{n+1} = \sum_{j=0}^{k-1} \nabla^j y'_n + \rho\,\nabla^k y'_n$$

$$y_{n+1} = p_{n+1} + h\,[(1-\rho)\gamma_k + \rho\,\gamma_{k+1}]\,(1+\alpha)\,[f(t_{n+1}, p_{n+1}) - p'_{n+1}]$$

$$y'_{n+1} = f(t_{n+1}, y_{n+1}),$$

where $\rho = \alpha = 0$, corresponds to the usual procedure. For $y' = \lambda y$ this leads to a polynomial of degree 2 in $h\lambda$, and degree $k+1$ in r. Let r_p and r_e denote the principal and the extraneous root of largest magnitude as before. Consider

$$E = \frac{|1 - e^{-h\lambda} r_p(h\lambda, q, \alpha)|^2}{|h\lambda|^2}$$

$$R = |r_e(h\lambda, q, \alpha)|^2 - |r_p(h\lambda, q, \alpha)|^2.$$

where E, the square of the relative error per step, is to be minimized subject to $R \ge 0$. Suppose values of $h\lambda$, q, and α are such that $R = 0$. If this choice of α is to be optimal for this q and $h\lambda$, then there must be no solution to the system

$$\frac{\partial E}{\partial \alpha} \delta\alpha \; + \; \frac{\partial E}{\partial q} \delta q \; < \; 0$$

$$\frac{\partial R}{\partial \alpha} \delta\alpha \; + \; \frac{\partial R}{\partial q} \delta q \; < \; 0.$$

Setting the determinant of this system to 0, gives the condition for an optimal formula,

$$\frac{\partial E}{\partial \alpha} \frac{\partial R}{\partial q} \; - \; \frac{\partial E}{\partial q} \frac{\partial R}{\partial \alpha} \; = \; 0$$

$$R \; = \; 0.$$

I have written a code to solve this problem for Adams formulas of arbitrary order. This project has been set aside for now. But in trying to solve this system near q = 6, for given values of hλ, and unknowns α and q, there were either bugs in my code or it doesn't much matter how α is selected.

References

[1] Gertrude Blanch, On Modified Divided Differences I, Math. Comp. 8 (1954) 1-11.
[2] M.P. Calvo and J.M. Sanz-Serna, Variable steps for symplectic integrators (Report 1991/3, Departamento de Matemàtica Aplicada y Computación, Facultad de Ciencias, Universidad de Valladolid, Valladolid, Spain, April 1991).
[3] S.A. Gill, A process for the step-by-step integration of differential equations in an automatic digital computing machine, Proc. Cambridge Philos. Soc., 47 (1951) 96-108.
[4] L.W. Jackson, The Computation of Coefficients of Variable Step Adams Methods (University of Toronto, Dept. Comp. Science Technical Report No. 94, June 1976).
[5] L.W.Jackson and A.E.Sedgwick, Vandermonde Matrices, Co-ordinate Transformations, and Adams' Method (University of Alberta, Dept. Comp. Science Report No. TR77-5 September, 1977).
[6] F.T. Krogh, A Variable Step Variable Order Multistep Method for the Numerical Solution of Ordinary Differential Equations (TRW Systems Group, Redondo Beach, CA, Report No. 99900-6229-R000, May 1967). (An effort has been made to cite some of my more obscure work, for those who otherwise wouldn't know of it. These papers are available on request.)

[7] F.T. Krogh, The Numerical Integration of Stiff Differential Equations (TRW Systems Group, Redondo Beach, CA, Report No. 99900-6573-R000, March 1968)

[8] F.T. Krogh, An Integrator Design (Jet Propulsion Laboratory, Pasadena, CA, Technical Memorandum 33-479, May 1971).

[9] F.T. Krogh, Algorithms for Changing the Step Size, SIAM J. Numer. Anal. 10 (1973) 949-965.

[10] F.T. Krogh, Changing Stepsize in the Integration of Differential Equations Using Modified Divided Differences, Proceedings of the Conference on the Numerical Solution of Ordinary Differential Equations, October 1972, pp. 22-71. (Lecture Notes in Mathematics, Vol. 362, Springer-Verlag Berlin, 1974).

[11] F.T. Krogh, Summary of Test Results with Variants of a Variable Order Adams Method, pp. 277-281 of *Numerical Methods for Differential Systems*, by L. Lapidus and W.E. Schiesser (Academic Press, New York, 1976).

[12] F.T. Krogh, Notes on Partitioning in the Solution of Stiff Systems, Proceedings International Conference on Stiff Computation, Vol. II, April 12-14, 1982 Park City Utah. (Sponsored by the U.S. Air Force Office of Scientific Research.)

[13] F.T. Krogh, DIVA/SIVA, Chapter 14.1 of MATH77, Release 4.0 (Jet Propulsion Laboratory, Pasadena, CA, JPL D-1341, Rev. C, May 1992). (The entire MATH77 package is available for about $1000 from, COSMIC, The University of Georgia, 382 East Broad St., Athens, GA 30602.)

[14] F.T. Krogh, Design of a New General Purpose code, DIVI, for Solving Initial Value Problems in Ordinary Differential Equations (Jet Propulsion Laboratory Section 372, Internal Computing Memorandum 545, Pasadena, CA 91109, April 1992). (Available via anonymous ftp, at "math.jpl.nasa.gov" in "pub/divi".)

[15] F.T. Krogh, and Kris Stewart, Implementation of Variable Step BDF Methods for Stiff ODE's, in "Numerical Methods for Solving Stiff Initial Value Problems," Proceedings, Oberwolfach, 28.6 - 4.7.1981, Edited by Germund Dahlquist and Rolf Jeltsch, August 1981.

[16] F.T. Krogh, and Kris Stewart, Asymptotic ($h \to \infty$) Absolute Stability for BDFs Applied to Stiff Differential Equations, ACM Trans. on Math. Soft. 10 (1984) 45-57.

[17] L.F. Shampine, and M.K. Gordon, *Computer Solution of Ordinary Differential Equations, The Initial Value Problem* (W.H. Freeman and Co., San Francisco 1975).

[18] R.D. Skeel, Construction of Variable-Stepsize Multistep Formulas, Math. of Comp. 47 (1986) 503-510 and S45-S52.

[19] Kris Stewart, *Semi-Implicit Backward Differentiation Formulas* (PhD Dissertation, The University of New Mexico, Albuquerque, April 1987).